Lecture 21 **Math 20 Fall 2014, Dartmouth College**

Here's the textbook's favorite example of a Markov chain - weather in the Land of Oz. If it's raining today in the Land of Oz, tomorrow it will rain with probability $1/2$, it will be a nice day with probability $1/4$ or it will snow with probability $1/4$. If it's a nice day today, tomorrow will either rain or snow with equal probability. If it's snowing, it will snow with probability $1/2$, or either rain or it will be a nice day - each with probability $1/4$. The corresponding transition matrix is:

$$P = \begin{array}{c} R \\ N \\ S \end{array} \begin{bmatrix} 1/2 & 1/4 & 1/4 \\ 1/2 & 0 & 1/2 \\ 1/4 & 1/4 & 1/2 \end{bmatrix}$$

Which can also be illustrated with a random walk on this graph:

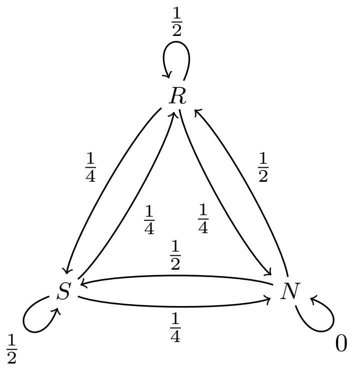

If it's a nice day today in the Land of Oz, what's the chance it will be a nice day the day after tomorrow?

Suppose we are walking on the graph on the left, starting at N. There are 3 different ways of being back at N 2 steps later - You can stay at N both times, go to R and back or go to S and back. Remember, when we talk about entries in a matrix, p_{ij} means "entry in row i and column j." In a transition matrix, p_{ij} is the probability that if the MC is in state i, it will be in state j on the next step. Notice how all the rows add up to 1?

If it's a nice day today, the probability that it will be a nice day the day after tomorrow is:

$$p_{NR}p_{RN} + p_{NN}p_{NN} + p_{NS}p_{SN} = \frac{1}{2}\frac{1}{4} + 0 + \frac{1}{2}\frac{1}{4} = \frac{1}{4}.$$

This can also be written as:

$$\sum_j p_{Nj}p_{jN}$$

And that's the (N, N) entry the matrix you get by squaring the matrix P. In fact, P^2 lists the probabilities of where the MC will be after TWO steps. The entry in row i, column j of P^2 is the probability that a MC starting at state i will be at state j after two steps.

Theorem 1 *Let P be a transition matrix of a Markov chain with states $\{S_1, \ldots, S_M\}$. The (i, j) entry of P^n is the probability that the Markov chain, starting in state S_i, will be in state S_j after n steps.*

You can find P, $P^2, \ldots,$ P^6 for the Weather in the Land of Oz in the textbook on page 408.

Exercise: It's Monday today in the Land of Oz. Can you tell me, without additional information, what the probability is that it will snow on Sunday? If it's Monday morning and looks like it will snow, rain or it will be a nice day each with probability $1/3$ today,what is the chance that it will snow on Thursday?

$$\begin{pmatrix} 1/3 & 1/3 & 1/3 \end{pmatrix} \begin{bmatrix} 1/2 & 1/4 & 1/4 \\ 1/2 & 0 & 1/2 \\ 1/4 & 1/4 & 1/2 \end{bmatrix} = \begin{pmatrix} 5/12 & 2/12 & 5/12 \end{pmatrix}$$

If instead we assumed that the probability distribution for weather on Monday will instead be $(0.4, 0.2, 0.4)$ for rain, nice day and snow respectively, the probabilities would e the same the next day:

$$\begin{pmatrix} 0.4 & 0.2 & 0.4 \end{pmatrix} \begin{bmatrix} 1/2 & 1/4 & 1/4 \\ 1/2 & 0 & 1/2 \\ 1/4 & 1/4 & 1/2 \end{bmatrix} = \begin{pmatrix} 0.4 & 0.2 & 0.4 \end{pmatrix}$$

we call this the stationary distribution - it's an eigenvector of P that has nonnegative entries that add up to 1. Every Markov chain has a stationary distribution, and often the stationary distribution is unique.

Definition 1 (Stationary Distribution) If P is an $m \times m$ transition matrix of a Markov chain, then the vector u such that all entries u_i are nonnegative, $\sum_{i=1}^{m} u_i = 1$ and:

$$uP = u,$$

is the *stationary distribution* of P.

Consider a simple random walk on:

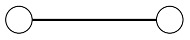

The corresponding transition matrix and its powers are:

$$P = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \qquad P^2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \qquad P^3 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \cdots$$

The stationary distribution of this walk is $(1/2, 1/2)$. Which means that if at the beginning we flip a coin to decide where the walk should start, then at any step the probability distribution for the current location is $(1/2, 1/2)$. But a particular walk, that has a given starting point, will not approach that distribution. Looking at the Land of Oz example gave us an impression that the overall probability distribution of the position of the walk should approach the stationary distribution. That is, in fact, often the case - but not in the walk above.

Theorem 2 *If some power P^i, $i \in \mathbb{N}$ of the transition matrix P has all entries positive, then:*

$$P(X_n = S_j) = u_j \text{ as } n \to \infty.$$

Equivalently,

$$P^n \to \begin{bmatrix} \underline{\qquad u \qquad} \\ \vdots \\ \overline{\qquad u \qquad} \end{bmatrix} \text{ as } n \to \infty.$$

This theorem doesn't apply to the simple random walk above, because all powers of P have 0 entries.

A Markov chain is called *irreducible* if for any starting state there is a way to get to any other state in a finite number of steps. Stationary distributions of irreducible Markov chains are unique. A simple random walk on this graph is an example of an MC that is not irreducible:

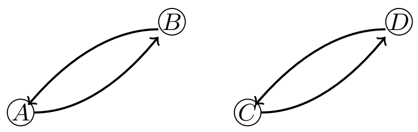

Can you find two different stationary distributions for this walk?

So, for graphs that do satisfy the condition of Theorem 2, the actual probability distribution of where the walk is at any given moment will tend to u. If you start the process somewhere, wait long enough, and freeze it, the current state will be distributed approximately according to u.